

ARTICLE

Toward Compliance Zero: AI and the Vanishing Costs of Regulatory Compliance

Paul Ohm

The Network Law Review is pleased to present a special issue entitled “The Law & Technology & Economics of AI.” This issue brings together multiple disciplines around a central question: What kind of governance does AI demand? A workshop with all the contributors took place on May 22–23, 2025, in Hong Kong, hosted by Adrian Kuenzler (HKU Law School), Thibault Schrepel (Vrije Universiteit Amsterdam), and Volker Stocker (Weizenbaum Institute). They also serve as the editors.

1. Toward Compliance Zero

Thanks to recent advances in artificial intelligence, we can now automate tasks that we have never before been able to automate.¹ Work that was expensive has become cheap;² tasks that once required a lot of time can now be accomplished nearly instantaneously.³ We can often now bypass the messy bottleneck of human labor.⁴ These developments have placed a bullseye on the back of knowledge and creative

¹ Christian Catalini, Jane Wu & Kevin Zhang, What Gets Measured, AI Will Automate, Harvard Business Review (June 19, 2025), <https://hbr.org/2025/06/what-gets-measured-ai-will-automate>.

² Id.

³ Shakked Noy & Whitney Zang, Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence, *Science* 381, 187 (2023).

⁴ Daron Acemoglu, Harms of AI, at 21 (NBER Working Paper No. 29247, 2021).

workers across industries, making us fear for the future prospects of lawyers, computer programmers, artists, authors, and educators.⁵

Another kind of worker hasn’t often been mentioned in these prognostications, yet there is no reason to omit them from the list of the disrupted: compliance experts.⁶ Any person who makes a living by helping companies, government agencies, and individuals comply with the burdens of government regulation will also soon be replaced by new forms of AI that can automate their work. These range from highly paid workers at name-brand firms like PricewaterhouseCoopers, McKinsey, or Covington, to workaday lawyers and consultants toiling in relative obscurity.⁷

I am not indulging in fantasies of Artificial General Intelligence (AGI). My prediction is based on an understanding of three factors: the work required for regulatory compliance, the AI advances that have been made to date, and modest and grounded predictions of what AI will soon be able to do in the near term. Compliance work entails precisely the kind of work that the latest advances in AI have disrupted. Compliance involves, among other things, interpreting complex regulatory texts, staying apprised of changes to the rules, and performing a series of knowledge-bound tasks such as assessing risk, summarizing documents, writing reports, and making disclosures. AI can assist or single-handedly accomplish all these tasks.⁸

If modest predictions of current and near-future AI capability come to pass, then AI automation will drive the cost of regulatory compliance down so close to zero that, going forward, we should consider most regulations to impose no compliance burden whatsoever.

This development will be a boon for investment and innovation. The era of crippling government compliance costs is over.

The automation of compliance will give rise to low-cost consultancies that use AI to specialize in helping companies comply with a broad set of government regulations. Cheap and ready access to these specialists will mean that small

⁵ Ali Zarifhonorvar, *Economics of ChatGPT: A Labor Market View on the Occupational Impact of Artificial Intelligence*, 10 *J. Eur. Bioethics & Data Ethics* 100, 113 (2023), <https://doi.org/10.1108/JEBDE-10-2023-0021>.

⁶ For a recent example, see Chip Cutter, *AI Is Coming for the Consultants. Inside McKinsey, ‘This Is Existential.’*, *Wall Street Journal*, Aug. 2, 2025, <https://www.wsj.com/tech/ai/mckinsey-consulting-firms-ai-strategy-89fbf1be>.

⁷ Marie-Therese Sekwenz, Rita Gsenger, Volker Stocker, Esther Görnemann, Dinara Talypova, Simon Parkin, Lea Greminger & Georgios Smaragdakis, *Can’t LLMs do that? Supporting Third-Party Audits under the DSA: Exploring Large Language Models for Systemic Risk Evaluation of the Digital Services Act in an Interdisciplinary Setting*, 4 *ACM Proc. Symposium on Human-Computer Interaction for Work (CHIWORK ’25 Adjunct)*, 1–12 (2025). <https://doi.org/10.1145/3707640.3731929>.

⁸ *Id.*

businesses and startups, not only deep pocketed giants, will enjoy the benefits of near-zero-cost compliance.

This will also be a boon for those promulgating, defending, and supporting new and aggressive forms of regulation. It will remove one of the most significant objections that can be lodged at a regulation: the high cost of compliance. For as long as we have had regulations, critics have assailed them for the way they create waste and inefficiency.⁹ Going forward, we can dismiss arguments like these as false and outdated, a relic of a bygone era. Regulatory compliance is now basically free.

2. The Road to Compliance Zero

In this short essay, I cannot offer a complete proof that we are heading for compliance zero. Instead, I present significant supporting evidence, confident that future work will more fully prove the thesis. I offer three forms of evidence. First, the compliance zero thesis is a simple extension of broader contemporary claims that have been made about AI's looming impact on knowledge work. Second, I survey the compliance obligations of the EU AI Act, demonstrating how they require mostly the kind of work that can be automated by large language models and agentic AI. Third, I point to a recent judicial opinion from a court in Hamburg, Germany, which provides one concrete example of how large language models can now automate a once burdensome compliance task, namely recognizing the “no scraping allowed” opt-out wishes of website owners.

2.1. AI and Knowledge Work

It is very important to clarify what I mean by “compliance burden.” I am talking about the effort and work required to comply with a law or regulation. I mean what might also be called paperwork, office work, or due diligence. But this goes beyond rote or mechanistic paperwork—I also mean sophisticated and complex steps that have, until now, required human interpretation, judgment, coordination, and organization.

This definition of compliance burden excludes the changes in behavior and system and organizational design mandated, encouraged, or incentivized by a law or regulation. In other words, it does not cover a law or rule's direct regulatory mechanisms. For example, a ban on facial recognition technology (FRT) may impose an expensive burden on police departments to dismantle preexisting FRT systems and to find alternative means of conducting investigations, but these are not the kinds of paperwork burdens I mean.

⁹ Francesco Trebbi & Miao Ben Zhang, *The Cost of Regulatory Compliance in the United States*, NBER Working Paper No. 30691, at 13–17, 27–28 (Nov. 2022), <http://www.nber.org/papers/w30691>.

Begin with two propositions: (1) most compliance work counts as legal work; and (2) compliance tasks tend to be on the simpler, more rote, less creative end of the spectrum of types of tasks that lawyers perform. Think simple contract review versus complex strategic advising.

Consider evidence to support these propositions. Several U.S.-based law schools now offer degree programs for compliance.¹⁰ Compared to a full JD program, these tend to be shorter (one year versus three), are often held exclusively online, require courses drawn from a limited subset of the law school curriculum, and are targeted at students who do not need to pass the bar or practice law.¹¹

If these propositions are mostly true, meaning compliance work tends to be a simplified subset of legal work, then every argument that has been made in the past few years about how AI will begin to automate legal work will have even greater force when it comes to compliance work.

The emerging literature debating and measuring whether AI will replace lawyers is already vast and rapidly expanding.¹² Controlled studies suggest that large language models can pass the bar exam¹³ and ace some law school exams.¹⁴ Researchers have conducted randomized controlled experiments in which law students have completed typical legal work both with and without using AI tools, and they suggest that these tools speed up legal tasks, make legal workers more productive, and produce higher quality work on most tasks.¹⁵

¹⁰ Online Master’s in Compliance, Fordham Law, <https://www.fordham.edu/school-of-law/admissions/msl-admissions/online-msl-in-compliance/> (last visited July 28, 2025); Online Master’s in Compliance With a Corporate Focus (MSL), Pitt Law, <https://online.law.pitt.edu/msl/corporate-compliance/> (last visited July 28, 2025); Master of Studies in Law (MSL) – Business Law & Compliance, Wake Forest Law, <https://law.wfu.edu/msl-blc/>; <https://www.colorado.edu/law/academics/degrees/msl-ethics-and-compliance> (last visited July 28, 2025).

¹¹ M.S.L. Admissions: Online Master’s in Compliance, Fordham Law, <https://www.fordham.edu/school-of-law/admissions/msl-admissions/online-msl-in-compliance/> (online-only, twelve-month degree program); Pitt Law’s Master of Studies in Law Curriculum, Pitt Law, <https://online.law.pitt.edu/msl/curriculum/> (online-only with set courses).

¹² See, e.g., Daniel Schwarcz et al., *AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice*, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5162111; Richard Susskind, *Tomorrow’s Lawyers: An Introduction to Your Future* (2023); Jonathan H. Choi, Amy Monahan, & Daniel Schwarcz, *Lawyering in the Age of Artificial Intelligence*, 109 Minn. L. Rev. 147 (2024); Aksh Garg & Megan Ma, *Opportunities and Challenges in Legal AI*, Stanford Law School (Jan. 2025); John G. Roberts, Jr., 2023 Year-End Report on the Federal Judiciary, Sup. Ct. of the U.S. 5 (2023).

¹³ Daniel Martin Katz, Michael James Bommarito, Shang Gao & Pablo Arredondo, *GPT-4 Passes the Bar Exam*, Phil. Trans. R. Soc. (Apr. 15, 2024), at 1, 3–5; see also Eric Martinez, *Re-Evaluating GPT-4’s Bar Exam Performance*, Artificial Intelligence and Law (2024).

¹⁴ 15] Andrew Blair-Stanek et al., *AI Gets Its First Law School A+s*, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5274547 (2025).

¹⁵ E.g., Schwarcz, *AI-Powered*, supra note 13, Part III.

The conventional wisdom is that AI will replace humans for some legal work, although many hold out hope that work requiring judgment and subtlety and creativity cannot be automated, at least not given the state of the art of AI.¹⁶ Compliance work tends to be on the automatable end of this spectrum.

Others have noted the likelihood that AI can and will be used for compliance tasks.¹⁷ Many vendors pitch their products and services specifically for regulatory compliance.¹⁸ Some use the label, RegTech, particularly for compliance in the FinTech industry.¹⁹ Recently, reports suggest that Meta will begin to automate “up to 90% of all risk assessments.”²⁰

2.2. Compliance Under the EU AI Act

To make this argument more concrete, consider a recently enacted law that has been decried as crippling burdensome:²¹ the European Union’s AI Act.²² To be clear, my claim extends beyond rules that regulate AI, and I could have illustrated the point with many other types of law, such as data protection, consumer protection, securities, or tax law. One reason to highlight an AI regulation is because the companies covered by these rules will also be the companies best positioned to harness the power of AI to reduce their regulatory burden.

¹⁶ Josef Valvoda et al., *The Ethics of Automating Legal Actors*, <https://arxiv.org/pdf/2312.00584>; Sayash Kapoor, Peter Henderson & Arvind Narayanan, *Promises and Pitfalls of Artificial Intelligence for Legal Applications*, <https://arxiv.org/pdf/2402.01656>.

¹⁷ Narendra Sharma, *AI-Powered Compliance: The New Frontier of Regulatory Innovation*, Union Leader (May 18, 2025), <https://www.unionleader.com/.../ask-the-expert-ai-powered-compliance-the-new-frontier-of-regulatory-innovation>; Rabihah Butler, *Generative AI and its Impact on Corporate Compliance*, Thomson Reuters (July 19, 2023), <https://www.thomsonreuters.com/.../generative-ai-corporate-compliance/>; Matt Kunkel, *Harnessing AI For Future-Proofing Regulatory Compliance*, Forbes (Jan. 25, 2024), <https://www.forbes.com/.../harnessing-ai-for-future-proofing-regulatory-compliance/>.

¹⁸ E.g., *Compliance.AI*, <https://www.compliance.ai/> (last visited July 25, 2025); Moritz Homann, *How Artificial Intelligence is Revolutionizing Compliance Management*, EQS (July 11, 2024), <https://www.eqs.com/.../how-artificial-intelligence-is-revolutionizing-compliance-management/>; Zvi Korn, *How AI Can Help Regulatory Compliance*, Datarails (July 12, 2024), <https://www.datarails.com/how-ai-can-help-regulatory-compliance/>.

¹⁹ Matt Morgan, *What Is Regulatory Technology, and How Are Businesses Using It?*, BizTech (Oct. 7, 2024), <https://biztechmagazine.com/article/2023/12/what-regulatory-technology-and-how-are-businesses-using-it-perfcon>.

²⁰ NPR, *Meta says it will automate up to 90% of risk assessments* (May 31, 2025), <https://www.npr.org/2025/05/31/nx-s1-5407870/meta-ai-facebook-instagram-risks>.

²¹ Elizabeth M. Renieris, David Kiron & Steven Mills, *Organizations Face Challenges in Timely Compliance With the EU AI Act*, MIT Sloan Management Review (June 13, 2024), <https://sloanreview.mit.edu/.../organizations-face-challenges-in-timely-compliance-with-the-eu-ai-act/>.

²² Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), 2024 O.J. (L 2024/1689) 1.

Every single supposedly burdensome compliance obligation or requirement introduced by the AI Act can be partially or entirely automated given recent advances in frontier models.

To start, the EU AI Act is a long and complex text. It is as long as a typical novel.²³ But a novel this isn't, and the language is densely written and jargon laden. Reading it once through would take a typical lawyer hours. Comprehending and organizing the Act's many interlocking requirements would require weeks of study by teams of lawyers. Modern LLMs excel at summarizing long texts. The entire EU AI Act can fit within most cutting-edge LLMs' "context windows," and produce a passable and accurate summary of the Act's obligations in seconds.²⁴ To be fair, some provisions of the AI Act have been criticized as vague, contradictory, or ambiguous,²⁵ problems that an LLM cannot alone remedy.

Moving beyond simply understanding the rules, the central focus of the AI Act is on technical documentation, record keeping, and transparency.²⁶ For systems deemed "high risk,"²⁷ the Act obligates companies to track dozens of categories of information, such as the "intended purpose" of the system, "the forms in which the AI system is placed on the market or put into service," "the methods and steps performed for the development of the AI system," "the design specification of the system," and "the description of the system architecture."²⁸ Large language models excel in analyzing texts, internally representing the subtle meaning of those texts, and summarizing those texts in human-understandable prose.²⁹

The AI Act obligates companies to analyze and summarize large datasets, for example training datasets.³⁰ Once again for high-risk systems, it requires "data governance and management practices" concerning "design choices," "data collection processes," "an assessment of the availability, quantity, and suitability of the datasets that are needed," and a number of other assessments necessary to assess

²³ The 113 operative Articles (including footnotes) contain more than 49,000 words. The thirteen annexes contain almost 8,000 words. The 180 recitals contain more than 35,000 words.

²⁴ What is a Context Window?, IBM (Nov. 7, 2024), <https://www.ibm.com/think/topics/context-window>.

²⁵ Alexander Procter, EU's AI Rules Are Coming And Nobody's Sure What They Mean, Spark Expert Insights by Okoone (Aug. 13, 2025), <https://www.okoone.com/.../eus-ai-rules-are-coming-and-nobodys-sure-what-they-mean/>.

²⁶ ISACA, Understanding the EU AI Act: Requirements and Next Steps (Oct. 18, 2024), <https://www.isaca.org/resources/white-papers/2024/understanding-the-eu-ai-act>.

²⁷ Regulation (EU) 2024/1689, Art. 6 (defining high-risk systems).

²⁸ Regulation (EU) 2024/1689, Annex IV.

²⁹ Vicente J. Bermejo et al., LLMs Outperform Outsourced Human Coders on Complex Textual Analysis, SSRN (Nov. 13, 2024), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5020034.

³⁰ ISACA, *supra* note 27 (describing the "data and data governance" obligations of the Act).

bias in the system.³¹ LLMs are also quite capable of automating the analysis of large amounts of data.³²

Many obligations in the AI Act require the production of long and detailed texts: assessments, risk summaries, compliance summaries, etc.³³ Providers of high-risk systems must generate “technical documentation” to permit the government to “assess the compliance” of their system under the Act.³⁴ They must also produce “instructions for use” of the system by deployers by conveying “concise, complete, correct, and clear information that is relevant, accessible, and comprehensible to deployers.”³⁵ LLMs are famously capable at writing coherent texts quickly.³⁶

Some AI Act requirements require the generation of images and videos. For example, the AI Act obligates providers and deployers of AI systems “to ensure, to their best extent, a sufficient level of ‘AI literacy’ of their staff and other persons” dealing with their AI systems.³⁷ One way to comply is by delivering tailored training and education materials for employees.³⁸ A European Commission FAQ opines that this might require more than simple instructions for use, in favor of “trainings and guidance.”³⁹ Many companies have produced video tutorials for their staff to comply with this requirement.⁴⁰ Image and video generation models are increasingly capable at generating visuals based on user inputs and prompts.⁴¹

The AI Act will require the evaluation and creation of computer source code and configuration code. For example, provisions of the law require cybersecurity assessments, remediation, and countermeasures.⁴² This obligates companies to put in place systems “to prevent, detect, respond to, resolve and control for attacks” in their systems.⁴³ LLMs are quite proficient at interpreting, modifying, and writing this kind of technical code.⁴⁴ Of course, cybersecurity requires far more than

³¹ Regulation (EU) 2024/1689, Art. 10.

³² Immanuel Trummer, *Data Analysis with LLMs* (2025).

³³ ISACA, *supra* note 27 (technical documentation obligations).

³⁴ Regulation (EU) 2024/1689, Art. 11.

³⁵ Regulation (EU) 2024/1689, Art. 13.

³⁶ Tom B. Brown et al., *Language Models are Few-Shot Learners*, arXiv (revised July 22, 2020), <https://arxiv.org/abs/2005.14165>.

³⁷ Regulation (EU) 2024/1689, Art. 4.

³⁸ Regulation (EU) 2024/1689, Recital 56.

³⁹ European Commission, *AI Literacy – Questions and Answers*, <https://digital-strategy.ec.europa.eu/en/faqs/ai-literacy-questions-answers>.

⁴⁰ European Commission, *Living Repository to Foster Learning and Exchange on AILiteracy* (Apr. 2025 update), <https://digital-strategy.ec.europa.eu/.../living-repository-foster-learning-and-exchange-ai-literacy>.

⁴¹ OpenAI, *Sora (text-to-video)*, <https://openai.com/sora/> (last visited July 28, 2025).

⁴² Regulation (EU) 2024/1689, Art. 15.

⁴³ Regulation (EU) 2024/1689, Art. 15(5).

⁴⁴ Ze Sheng et al., *LLMs in Software Security: A Survey of Vulnerability Detection Techniques and Insights* (Feb. 12, 2025), <https://arxiv.org/pdf/2502.07049>.

detection, and it is one area in which LLM assistance will only be partial; we will still need human expertise to fully comply with parts of the AI Act, at least for now.

Other advances in AI will help stitch together these disparate tasks. So-called reasoning models excel at breaking down complex tasks into discrete separate pieces.⁴⁵ Mixture-of-experts approaches deploy different strategies to solving problems, helping find optimal solutions.⁴⁶ Increasingly advanced tools orchestrate multiple AI agents each performing different tasks, combining their inputs into a pipeline of work.⁴⁷

2.3. LAION v Robert Kneschke

In addition to this back-of-the-envelope analysis of the EU AI Act, consider a more discrete example drawn from real-world legal practice. A trial judge in Hamburg, Germany, recently demonstrated how the law will be interpreted to recognize the way AI models can be used to automate legal compliance in ways that would have felt like science fiction only a few years ago. Understanding the analysis requires a short primer on European Union and German copyright law.

The case involved the activity of a nonprofit organization called LAION.⁴⁸ LAION built a dataset that contained the URL locations of billions of images by analyzing the data crawled by another group, called Common Crawl.⁴⁹ Many of the images associated with LAION’s URLs were protected under German copyright law. LAION makes this list of URLs (but not the images themselves) free for all as an internet download,⁵⁰ and many image generation and multimodal models, such as Stable Diffusion, have used the LAION list to find and download millions of images.⁵¹ These images have served as crucial training data for these

⁴⁵ OpenAI, o1 System Card (Sept. 12, 2024), <https://cdn.openai.com/o1-system-card.pdf>.

⁴⁶ DeepSeek-AI, DeepSeek-V3 Technical Report (Feb. 18, 2025), <https://arxiv.org/pdf/2412.19437>.

⁴⁷ Lei Wang et al., A Survey on Large Language Model based Autonomous Agents (last revised Mar. 2, 2025), <https://arxiv.org/abs/2308.11432>.

⁴⁸ LG Hamburg [Regional Court Hamburg], Sept. 27, 2024, 310 O 227/23, available at Hellenic Copyright Organization, [https://opi.gr/en/news/5388/\(Ger.\)](https://opi.gr/en/news/5388/(Ger.)). The court did not make an English translation available; analysis relies on an unofficial translation: Chat GPT Is Eating the World, Unofficial English Translation (PDF).

⁴⁹ Christoph Schuhmann et al., LAION-5B: An Open Large-Scale Dataset for Training Next Generation Image-Text Models (Oct. 16, 2022), <https://arxiv.org/abs/2210.08402>.

⁵⁰ LAION, Projects, <https://laion.ai/projects/>.

⁵¹ E.g., Robin Rombach et al., High-Resolution Image Synthesis with Latent Diffusion Models (rev. Apr. 13, 2022), <https://arxiv.org/abs/2112.10752>.

models,⁵² and these downloads have been characterized as massive infringements of copyright.⁵³

A group of the owners of some of these copyrights sued LAION in German court for infringement.⁵⁴ Focus only on one narrow issue litigated in this case: a caveat to an exception to an exception to the ordinary prohibition against infringement.

The exception is the Text and Data Mining (TDM) exception.⁵⁵ It is a defense to copyright infringement to scrape images for the purpose of text and data mining.⁵⁶ LAION asserted that because its list is primarily useful for TDM, it can claim this exception.⁵⁷

The TDM exception has a significant exception to the exception, at least for TDM that cannot qualify as “scientific research purposes.”⁵⁸ That exception to the exception applies when the copyright owner invokes a “reservation of rights” alongside the copyrighted content.⁵⁹ In other words, it creates an “opt out” opportunity for copyright owners to the TDM exception.⁶⁰

Finally, the caveat: the copyright owner must convey this reservation of rights to the web scraper in a “machine readable” way.⁶¹ This phrase had already received considerable expert commentary and analysis.⁶² Much of this commentary focused on the so-called robots.txt protocol, a widely but informally recognized internet

⁵² Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk & Jenia Jitsev, LAION: Pioneering the democratization of AI research and bridging gaps in the field, *Falling Walls Science Summit* (2023), <https://apply.falling-walls.com/.../laion-pioneering-the-democratisation-of-ai-research-and-bridging-gaps-in-the-field/>.

⁵³ Chloe Xiang, A Photographer Tried to Get His Photos Removed from an AI Dataset. He Got an Invoice Instead, *VICE* (Apr. 28, 2023).

⁵⁴ LG Hamburg, 310 O 227/23 (Ger.).

⁵⁵ Directive 2019/790 of the European Parliament and of the Council of 17 Apr. 2019 on copyright and related rights in the Digital Single Market and amending Directives 96/9/EC and 2001/29/EC, 2019 O.J. (L 130) 92 (EU).

⁵⁶ Directive 2019/790, recitals 8–10.

⁵⁷ LG Hamburg, at 10; see also EUIPO, Recent Case Law: Germany—Hamburg District Court, 310 O 227/23, *LAION v. Robert Kneschke* (Sept. 27, 2024), <https://www.euipo.europa.eu/.../laion-v-robert-kneschke>.

⁵⁸ EU and German law contain two related TDM exceptions: Article 3 (scientific research) (transposed in German Law, § 44a UrhG) and Article 4 (other purposes) (transposed in German Law, § 60d UrhG). Only Article 4 allows rights reservation (Art. 4(3)).

⁵⁹ Directive 2019/790, recital 10, art. 3.

⁶⁰ Ana Lazarova et al., Statement on the Opt-Out Exception Regime / Rights-Reservation Regime for Text and Data Mining under Article 4 of the EU Directive on Copyright in the Digital Single Market, Creative Commons (Dec. 17, 2021), <https://creativecommons.org/.../CC-Statement-on-the-TDM-Exception-Art-4-DSM-Final.pdf>.

⁶¹ Directive 2019/790, recital 18.

⁶² Louise de Béthune, Balancing the Scales: Navigating Text and Data Mining, Awaiting Standardisation, *CiTIP Blog* (KU Leuven) (Mar. 18, 2025), <https://www.law.kuleuven.be/.../balancing-the-scales-navigating-text-and-data-mining-awaiting-standardisation/>.

standard.⁶³ According to the standard, the owner of a webpage may include a file called robots.txt on its web server and within the file, it can specify which web scraping bots are welcome or not.⁶⁴ Well-behaved (i.e., standard-compliant) scraping bots will respect the wishes expressed in a robots.txt file. The Internet Engineering Task Force (IETF), the standards body that promulgated the original robots.txt protocol, has actively been debating whether to modify it to better serve the TDM opt-out purpose.⁶⁵

Regardless, many of the images in the LAION dataset were not protected by a machine readable robots.txt entry.⁶⁶ In the Hamburg litigation, the copyright owners argued that they had expressed their reservations of rights, through “terms of service” documents that prohibited scraping for certain purposes, including for training machine learning models.⁶⁷ Although terms of service are written in human language (in many cases, English) and ostensibly published for human comprehension, these plaintiffs argued that these documents were sufficiently machine readable to qualify as a TDM reservation of rights under German copyright law.⁶⁸

The Hamburg trial court agreed, expressing the kind of logic that is at the core of this essay.⁶⁹ “Machine readable,” within the meaning of the act, includes anything that is possible using “state-of-the-art technologies,” the court said.⁷⁰ Today’s state-of-the-art includes Large Language Models (LLMs) embedded in web scraping systems, tools that seem amply capable of noticing the presence of a terms-of-service hyperlink, following that link to download the terms, finding the part of the terms that address web scraping, and accurately interpreting whether scraping for TDM purposes has been forbidden.⁷¹

⁶³ Shayne Longpre et al., Consent in Crisis: The Rapid Decline of the AI Data Commons (rev. July 24, 2024), <https://arxiv.org/abs/2407.14933>.

⁶⁴ IETF, A File Format to Aid in Accessibility of Web Content for Automated Processing (RFC 9309) (Aug. 2022), <https://datatracker.ietf.org/doc/html/rfc9309>; Justin Hendrix, Robots.txt Is Having a Moment. Here’s Why We Should Care, Tech Policy Press (Apr. 9, 2024), <https://www.techpolicy.press/.../robotstxt-is-having-a-moment-heres-why-we-should-care/>.

⁶⁵ IAB Workshop on AI-CONTROL (aicontrolws), IETF, <https://datatracker.ietf.org/group/aicontrolws/about/> (last visited July 28, 2025).

⁶⁶ LAION FAQ: <https://laion.ai/faq/> (“Common Crawl did the crawling and respected robots.txt. We analyze their data and evaluate images.”).

⁶⁷ Adrian Aronsson-Storrier & Oliver Fairhurst, How to ‘Opt-Out’ Your Images from Use in AI Training: German Court Considers the EU Text and Data Mining Copyright Exceptions, Lewis Silkin (Oct. 7, 2024), <https://www.lewissilkin.com/.../how-to-opt-out-your-images-from-use-in-ai-training>.

⁶⁸ Morrison & Foerster, To Scrape or Not to Scrape? First Court Decision on the EU Copyright Exception for Text and Data Mining in Germany (Oct. 4, 2024), <https://www.mofo.com/.../241004-to-scrape-or-not-to-scrape-first-court-decision>.

⁶⁹ LG Hamburg at 12–14 (English translation).

⁷⁰ Id. at 13.

⁷¹ Id.

The court found that this kind of ability was well within the capabilities of recent advances in LLMs.⁷² It noted the special irony of these particular defendants disclaiming this power, hoisting them on their own AI petard:⁷³ “it would, in the Chamber’s view, present a certain inconsistency in valuation to allow the development of increasingly powerful text-understanding and text-generating AI models through the [TDM] exception . . . , while simultaneously not requiring the use of already existing AI models within the [reservation of rights provision.]”⁷⁴

2.4. To the AI Skeptics

Some readers will still be skeptical. For every believer in the unbounded potential of AI power, there is an equally confident skeptic convinced that such claims are overinflated hype. The skeptics point to the persistent tendency of LLMs to hallucinate, the brittleness of guardrails, and the track record of the tech industry to overpromise and underdeliver to support their disbelief. LLMs are not ready for prime time, these people argue, not even for relatively constrained tasks like regulatory compliance.

I am not so skeptical. I have been convinced that LLMs are quite capable at performing complex tasks that could not have been automated a decade ago, and that they will begin to replace the human labor used for many knowledge industry tasks.

But if the skeptics are right—if AI solutions are not nearly as powerful as some claim them to be, and if AI cannot capably perform compliance work, then time will tell. And if future developments reveal that the skeptics are right and LLMs are too error-prone and irredeemably hallucinatory to conduct somewhat rote compliance work, even though it would disprove this essay’s core thesis, it would nevertheless strengthen arguments in favor of the beefed-up regulatory compliance of the AI industry, which is a principal takeaway of this essay. For if AI is too error prone for even simple compliance work, we surely must regulate it to prevent the inevitable harms those errors will cause.

3. What This Means, and Doesn’t Mean

Advances in AI will tend to drive the costs of regulatory compliance toward zero beyond the regulation of AI. In every possible sphere of industrial regulation, many once labor-intensive, impossible-to-automate compliance tasks can now be performed by AI, and we will soon see a dramatic reduction in the amount of time and money spent on compliance. From tax to climate to consumer protection and

⁷² LG Hamburg at 13.

⁷³ The phrase “hoist on one’s own petard” originates with Shakespeare, *Hamlet*, act 3, sc. 4. See Melissa Mohr, *The Explosive Origin of ‘Hoist by One’s Own Petard’*, *Christian Science Monitor* (Sept. 27, 2021), <https://www.csmonitor.com/.../The-explosive-origin-of-hoist-by-one-s-own-petard>.

⁷⁴ LG Hamburg at 14.

beyond, the ability of LLMs to gather, process, interpret, and generate texts; the ability of image-and-video generation models to generate visuals; and the ability of AI agentic systems to perform complex multi-step data tasks will streamline the way companies comply with regulation.

To be clear, this does not mean that we can stop debating the efficacy or wisdom of new regulation. Quite the contrary, this will have the simple, salutary effect of removing one major objection to forms of new regulation: compliance costs. Saying that a regulation poses no significant costs of compliance is not the same thing as saying the law is a good law. We will be liberated to debate non-compliance-cost-based arguments in favor or against a new regulation, for example whether it addresses an important problem, whether it is sufficiently tailored to the problem, whether it improperly selects winners and losers, and whether it assigns a proper role to government. I predict that many proposed regulations will fall in this gantlet of considerations. The difference is we can avoid debating whether a proposal imposes an undue compliance burden: going forward, the answer will always be “no.”

We must also be mindful that the promise of compliance zero could lead us to expand the regulatory state in ways we should resist. If we are not judicious, regulating with compliance zero might tend to produce laws that are more complex and technocratic and beyond human comprehension than the laws we create when humans perform compliance. After all, the complexity of any legal regime can now be met by the complexity of AI-powered compliance responses. Within an increasingly byzantine call-and-response of regulation and compliance, we may create warrens of arbitrage, fraud, and unfair competition. The trick is to recognize the way that automation reduces the burden of compliance, without using it as an excuse to maximize complexity. In other words, the fact that companies can comply with any law with minimal burden should not liberate us from the goal of crafting human-scale, explainable, interpretable, frictionful regulation.

This short essay will not be the final word on these matters; there is much left for future work. First, although we are sliding down an asymptote toward compliance zero, we will never reach truly cost-free compliance. Setting up agentic workflows to do compliance work will cost something, and running a large system of AI agents imposes real costs, both financial and environmental. We need to learn to measure and accurately characterize these irreducible costs. Second, some regulations will impose compliance work that will indeed be difficult to automate, despite AI. For example, regulations that require humans to work with one another—say a rule that requires board scrutiny, review by an oversight committee, or consultation with users or the public—cannot be fully automated, by definition. Poorly drafted or ambiguous regulations will also resist automation. We should study the kinds of

compliance burdens that cannot be automated, to teach us both about regulation and about AI.

The burden of regulatory compliance will soon be nothing but a bad memory of days past, a story the old-timers tell the young ones about how the law used to operate. Thanks to advances in AI, most compliance work will soon be automated. Liberated from worrying about the costs imposed by our rules, our challenge is to write rules that are effective, tailored, and wise.

Paul Ohm⁷⁵

⁷⁵ Professor of Law, Georgetown University Law Center. Thanks to the organizers and participants of the Law, Technology, and Economics of AI Regulation workshop of the Platforms, Markets, and the Digital Society project. Special thanks to organizers Adrian Kuenzler, Thibault Schrepel, and Volker Stocker. This work benefited greatly from conversations with Steve Albrecht, William Lehr, Jason Potts, Michal Shur-Ofry, and Harry Surden. Finally, thanks to Taylor Courtney for excellent research assistance.